

AMENDMENTS TO THE CLAIMS

Please amend claims 1, 13, 28, and 29 and cancel claims 3 and 30-32 as shown below. Pursuant to 37 C.F.R. § 1.121(c), the text of all pending claims, along with their current status, is set forth below.

1. (Currently Amended) A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, writes to the first and second remote storage entities being tracked via respective first and second sidefiles, the first and second storage entities having different levels of write-currency relative to each other, the method comprising:

polling the second storage entity by the primary storage entity to determine an newest contiguous sequence number of the second sidefile;

generating acknowledgment information based on the newest contiguous sequence number of the second sidefile;

receiving acknowledgements from the second remote storage entity at both the primary storage entity and sending the acknowledgment information from the primary storage entity to the first remote storage entity;

discarding at least some of the sequence numbers from the first sidefile that are equal to or older than the newest contiguous sequence number, based on the acknowledgment information;

comparing acknowledgements and sequence numbers in the first sidefile with acknowledgements and sequence numbers in the second sidefile; and

updating writes stored at the second remote storage entity based upon the comparison of the first and second sidefiles.

2. (Original) The method of claim 1, wherein the updating of writes includes forwarding to the second remote storage entity writes that are present in the first sidefile but not yet present in the second sidefile.

3. (Canceled).

4 (Original) The method of claim 1, further comprising:

establishing a communication connection between the first and second remote storage entities in response to the primary storage entity becoming inoperative;

wherein the comparing and updating are also performed in response to the primary storage entity becoming inoperative.

5. (Previously Presented) The method of claim 1, further comprising:

configuring the first and second sidefiles to store a fixed number of writes therein, respectively; and

adding a new write to the first and second sidefiles by overwriting the oldest write therein, respectively.

6. (Previously Presented) The method of claim 1, further comprising:

adaptively adjusting the number of writes that can be stored in the first sidefile.

7. (Previously Presented) The method of claim 6, wherein the adaptive adjustment is based upon the writes that are stored in the second sidefile.

8. (Previously Presented) The method of claim 7, wherein the adaptive adjustment more particularly is based upon the write stored in the second sidefile which has the oldest contiguous sequence number.

9. (Previously Presented) The method of claim 7, further comprising:

identifying at least some of the writes stored in the second sidefile; and then accordingly informing the first remote storage entity regarding such identities.

10. (Previously Presented) The method of claim 1, further comprising:
 configuring the first sidefile to include a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.
11. (Previously Presented) The method of claim 1, further comprising:
 providing a third sidefile via which writes received thereby are tracked; and
 configuring the third sidefile to include a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.
12. (Previously Presented) The method of claim 11, further comprising:
 configuring the third sidefile further to include a field that is used to track whether a write-acknowledgement forwarded from the second remote storage entity has been acknowledged as having been received by the first remote storage entity.
13. (Currently Amended) A redundant data storage arrangement comprising:
 a primary storage entity configured to (1) forward writes to each of a mirroring first and second remote storage entity, the first and second storage entities having different levels of write-currency relative to each other and (2) forward acknowledgements from the second remote storage entity to the first remote storage entity;
 the mirroring first remote storage entity, ~~in communication with the primary storage entity,~~ which includes a first sidefile via which writes forwarded from the primary storage entity are tracked; and
 the mirroring second remote storage entity, ~~in communication with the primary storage entity,~~ which includes a second sidefile via which writes forwarded from the primary storage entity are tracked;
 ~~the first and second storage entities having different levels of write-currency relative to each other;~~

wherein the primary storage entity is configured to poll the mirroring second storage entity to determine a newest contiguous sequence number of the second sidefile, generate acknowledgment information based on the newest contiguous sequence number of the second sidefile, and send the acknowledgment information to the mirroring first remote storage entity;

wherein the mirroring first remote storage entity is configured to discard at least some of the sequence numbers from the first sidefile that are equal to or older than the newest contiguous sequence number, based on the acknowledgment information;

an initiating one of the first and second remote storage entities being operable to compare ~~acknowledgements and~~ sequence numbers in the first and second sidefiles[,] and invoke an updating of writes stored at the second remote storage entity based upon the comparison ~~of the first and second sidefiles~~.

14. (Previously Presented) The redundant data storage arrangement of claim 13, wherein the updating is performed by the first remote storage entity, which is operable to do so by forwarding to the second remote storage entity writes that are present in the first sidefile but not yet present in the second sidefile.

15. (Previously Presented) The redundant data storage arrangement of claim 13, wherein each of the first and second remote storage entities is operable to

preserve in the respective sidefile sequence numbers associated with the writes;
and

sort the respective sidefile according to the sequence numbers.

16. (Previously Presented) The redundant data storage arrangement of claim 13,
wherein:

the initiating one is further operable to establish a communication connection between itself and the other remote storage entity in response to the primary storage entity becoming inoperative; and

the comparison and the update are performed in response to the primary storage entity becoming inoperative.

17. (Previously Presented) The redundant data storage arrangement of claim 13, wherein the first remote storage entity is closer in proximity to the primary storage entity than the second remote storage entity.

18. (Previously Presented) The redundant data storage arrangement of claim 13, wherein:

the first and second sidefiles are configured to store a fixed number of writes therein, respectively; and

the first and second remote storage entities are operable to add a new write to the first and second sidefiles by overwriting the oldest write therein, respectively.

19. (Previously Presented) The redundant data storage arrangement of claim 13, wherein:

the first remote storage entity is operable to adaptively adjust the number of writes that can be stored in the first sidefile.

20. (Previously Presented) The redundant data storage arrangement of claim 19, wherein the adaptive adjustment is based upon the writes are stored in the second sidefile.

21. (Previously Presented) The redundant data storage arrangement of claim 20, wherein the adaptive adjustment more particularly is based upon the write stored in the second sidefile which has the newest contiguous sequence number.

22. (Previously Presented) The redundant data storage arrangement of claim 20, wherein:

the primary storage entity is operable to identify at least some of the writes stored in the second sidefile and then accordingly inform the first remote storage entity regarding such identities.

23. (Previously Presented) The redundant data storage arrangement of claim 13, wherein:

the first sidefile includes a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

24. (Previously Presented) The redundant data storage arrangement of claim 13, wherein:

the primary storage entity includes a third sidefile to track writes received thereby;
the first sidefile including a field that is used to track whether a write has been acknowledged by the second remote storage entity as having been received.

25. (Previously Presented) The redundant data storage arrangement of claim 24, wherein the first sidefile further includes a field that is used to track whether a write-acknowledgement forwarded from the second remote storage entity has been acknowledged as having been received by the first remote storage entity.

26. (Previously Presented) The redundant data storage arrangement of claim 13, wherein each of the first and second remote storage entities represents a tracked write in the first and second sidefile, respectively, with: location information as to where on a physical medium the write is to be performed; actual data associated with the write that is to be written to the physical medium; and a sequence number uniquely associated with the write.

27. (Previously Presented) The redundant data storage arrangement of claim 13, wherein:

the first remote storage entity receives writes forwarded synchronously from the primary storage entity; and

the second remote storage entity receives writes forwarded asynchronously from the primary storage entity.

28. (Currently Amended) A data synchronization method for a redundant data storage arrangement in which there are at least a primary storage entity and mirroring first and second remote storage entities in communication therewith, respectively, the method comprising:

synchronously receiving writes at the first remote storage entity that have been forwarded from the primary storage entity;

maintaining a first sidefile via which are tracked items that include sequence numbers assigned to writes, writes received at the first remote storage entity, and writes received at the second remote storage entity;

~~acknowledgements from the second remote storage entity received at both the primary storage entity and the first remote storage entity;~~

~~sequence numbers assigned to writes;~~

~~writes received at the first remote storage entity; and~~

~~writes received at the second remote storage entity; and~~

receiving acknowledgements from the primary storage entity, the acknowledgements indicating a newest contiguous sequence number of a second sidefile maintained by the mirroring second remote storage entity,

discarding at least some of the sequence numbers from the first sidefile that are equal to or older than the newest contiguous sequence number of the second side file; and

comparing ~~both acknowledgements and~~ the sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.

29. (Currently Amended) A mirroring first remote storage entity for a redundant data storage arrangement in which there are at least the first remote storage entity and a primary storage entity in communication therewith, and a mirroring second remote storage entity in communication with the primary storage entity, respectively, the first remote storage entity comprising:

memory to store data; and

a processor operable to:

synchronously receive writes that have been forwarded from the primary storage entity;

maintain a first sidefile in the memory via which are tracked items that include sequence numbers assigned to writes, writes received at the first remote storage entity, and writes received at the second remote storage entity;

~~acknowledgements from the second remote storage entity received at both the primary storage entity and the first remote storage entity;~~
~~sequence numbers assigned to writes,~~
~~writes received at the first remote storage entity, and~~
~~writes received at the second remote storage entity; and~~

receive acknowledgements from the second remote storage entity, the acknowledgements indicating a newest contiguous sequence number of a second sidefile maintained by the mirroring second remote storage entity,

discard at least some of the sequence numbers from the first sidefile that are equal to or older than the newest contiguous sequence number of the second side file; and

compare ~~both acknowledgements and~~ the sequence numbers at the first remote storage entity to determine whether to update writes stored at the second remote storage entity after a failure of the primary storage entity.

30-33. (Canceled).